

IN THE CLAIMS:

Please write the claims to read as follows:

- 1 1. (Original) A storage system for use in a storage system cluster, the storage system
2 comprising:
3 a storage operating system including a cluster connection manager adapted to cre-
4 ate, destroy, and maintain one or more communication sessions with a cluster partner, the
5 cluster connection manager operatively interconnected with a set of cluster connection
6 manager clients.
- 1 2. (Original) The storage system of claim 1 wherein one of the set of communication
2 clients comprises a failover monitor.
- 1 3. (Original) The storage system of claim 1 wherein one of the set of cluster connection
2 manager clients comprises a non-volatile random access memory shadowing process.
- 1 4. (Original) The storage system of claim 1 wherein the cluster connection manager is
2 further adapted to perform connection management operations in response to communi-
3 cations from the connection manager clients.
- 1 5. (Original) The storage system of claim 4 wherein the communications comprise an ap-
2 plication program interface function call.

1 6. (Original) The storage system of claim 1 wherein the cluster connection manager is
2 further adapted to load balance the one or more communication sessions over a plurality
3 of cluster interconnect devices,

1 7. (Original) The storage system of claim 1 wherein the cluster connection manager is
2 further adapted to perform a failover procedure for one or more communication sessions
3 from a failed cluster interconnect device to an operational cluster interconnect device.

1 8. (Original) The storage system of claim 1 wherein the cluster connection manager is
2 operatively interconnected with a plurality of cluster interconnect devices.

1 9. (Original) The storage system of claim 1 wherein the storage operating system com-
2 prises a plurality of cluster connection managers.

1 10. (Original) A storage operating system, executing on a storage system, the storage op-
2 erating system comprising:

3 a cluster connection manager adapted to manage a set of peer-to-peer connections
4 associated with a set of cluster connection manager clients executing on the storage sys-
5 tem.

1 11. (Original) The storage operating system of claim 10 wherein the set of cluster con-
2 nection manager clients comprises a failover monitor.

1 12. (Original) The storage operating system of claim 10 wherein the cluster connection
2 manager is further adapted to perform load balancing of the set of peer-to-peer connec-
3 tions over a plurality of cluster interconnect devices.

1 13. (Original) The storage operating system of claim 10 wherein the cluster connection
2 manager is further adapted to failover the set of peer-to-peer connections from a failed
3 cluster interconnect device to an operational cluster interconnect device.

1 14. (Currently Amended) A method for initiating a peer-to-peer communication session,
2 the method comprising the steps of:

3 creating, using a cluster connection manager executing on a storage system, an
4 initial connection with a cluster partner;
5 exchanging a set of peer connection information;
6 passing a set of client information to the cluster partner;
7 creating a set of appropriate communication ports;
8 alerting the cluster partner of a ready status; and
9 alerting a set of clients that the cluster partner is in a ready state.

1 15. (Original) The method of claim 14 wherein the set of clients comprises a failover
2 monitor process.

1 16. (Original) The method of claim 14 wherein the set of peer connection information
2 comprises a version number.

1 17. (Original) The method of claim 14 wherein the step of passing a set of client informa-
2 tion to the cluster partner further comprises the steps of:

3 collecting, from a set of clients, the set of client information; and

4 transferring the collected set of client information to the cluster

1 18. (Original) The method of claim 17 wherein the client information comprises a num-
2 ber of communication ports required.

1 19. (Original) The method of claim 17 wherein the set of client information further com-
2 prises an amount of memory requested by a particular client.

1 20. (Original) The method of claim 14 wherein the step of creating an initial connection
2 further comprises the step of using remote direct memory access primitives to create the
3 initial connection.

1 21. (Original) The method of claim 14 wherein the step of creating an initial connection
2 further comprises the step of performing a series of remote direct memory access opera-
3 tions to create the initial connection.

1 22. (Currently Amended) A method for terminating a peer-to-peer communication ses-
2 sion, the method comprising the steps of:

3 alerting, using a cluster connection manager executing on a storage system, a set
4 of clients of an impending termination of the communication session;

5 closing, by the clients, a set of communication ports associated with the commu-
6 nication session; and

7 performing an initialization of a peer-to-peer communication session procedure.

1 23. (Original) The method of claim 22 wherein the set of communication ports comprises
2 a set of virtual interface connections.

1 24. (Original) The method of claim 22 wherein the set of clients comprises a failover
2 monitor.

1 25. (Original) A storage operating system, executing on a storage system, the storage op-
2 erating system comprising:

3 a cluster connection manager having means to manage a set of peer-to-peer con-
4 nections associated with a set of cluster connection manager clients executing on the
5 storage system.

1 26. (Original) The storage operating system of claim 25 wherein the set of cluster con-
2 nection manager clients further comprises a failover monitor.

1 27. (Original) The storage operating system of claim 25 wherein the set of cluster con-
2 nection manager clients further comprises a nonvolatile random access memory shadow-
3 ing process.

1 28. (Original) A system configured to manage reliable peer communication among stor-
2 age systems in a clustered environment, the system comprising:

3 one or more peer processes executing on each storage system partner; and

4 a cluster connection manager executing on each storage system partner, the clus-
5 ter connection manager creating a set of peer-to-peer connections between the one or
6 more peer processes executing on each storage system.

Please add the following new claims 29 *et seq.*:

1 29. (New) A computer readable medium for initiating a peer-to-peer communication ses-
2 sion, the computer readable medium including program instructions for performing the
3 steps of:

4 creating, using a cluster connection manager executing on a storage system, an
5 initial connection with a cluster partner;
6 exchanging a set of peer connection information;
7 passing a set of client information to the cluster partner;
8 creating a set of appropriate communication ports;
9 alerting the cluster partner of a ready status; and
10 alerting a set of clients that the cluster partner is in a ready state.

1 30. (New) A computer readable medium for terminating a peer-to-peer communication
2 session, the computer readable medium including program instructions for performing
3 the steps of:

4 alerting, using a cluster connection manager executing on a storage system, a set
5 of clients of an impending termination of the communication session;
6 closing, by the clients, a set of communication ports associated with the commu-
7 nication session; and
8 performing an initialization of a peer-to-peer communication session procedure.

1 31. (New) A method for maintaining a peer-to peer communication, the method compris-
2 ing:

3 waiting for an event from a client communicating with a cluster partner to be re-
4 ceived by a cluster connection manager executing on a storage operating system;
5 determining whether the event is a client event; and
6 in response to determining that the event is a client event, performing the event
7 utilizing the cluster connection manager.

1 32. (New) The method of claim 31, further comprising:
2 in response to determining that the event was not a client event, alerting a set of
3 clients of an impending termination of the communication session;
4 closing, by the clients, a set of communication ports associated with the commu-
5 nication session; and
6 performing an initialization of a peer-to-peer communication session procedure.

1 33. (New) The method of claim 32 wherein the set of communication ports comprises a
2 set of virtual interface connections.

1 34. (New) The method of claim 32 wherein the set of clients comprises a failover moni-
2 tor.

1 35. (New) The method of claim 31 further comprising monitoring the status of one or
2 more cluster interconnect drivers utilizing the cluster connection manager.

1 36. (New) A computer readable medium for maintaining a peer-to-peer communication
2 session, the computer readable medium including program instructions for performing
3 the steps of:

4 waiting for an event from a client involved in a communication session to be re-
5 ceived by a cluster connection manager executing on a storage operating system;
6 determining that the event is a client event; and
7 in response, performing the event utilizing the cluster connection manager.

1 37. (New) A storage operating system, executing on a storage system, the storage operat-
2 ing system comprising:
3 one or more peer processes executing on each storage system partner;
4 a plurality of cluster interconnect drivers executing on the storage system; and
5 one or more cluster connection managers configured to detect a failure of a first
6 cluster interconnect driver and in response to determining the failure of the first cluster
7 interconnect driver, utilize a second cluster interconnect driver to access each storage sys-
8 tem partner.

1 38. (New) A storage operating system, executing on a storage system, the storage operat-
2 ing system comprising:
3 one or more peer processes executing on each storage system partner;
4 a plurality of cluster interconnect drivers executing on the storage system; and
5 one or more cluster connection managers configured to detect a high bandwidth
6 load on a first cluster connection manager and in response to detecting a high band width
7 load, utilize a second cluster connection manager to access each storage system partner.